

Attributing Fair Decisions with Attention Interventions

Ninareh Mehrabi, Umang Gupta, Fred Morstatter, Greg Ver Steeg, Aram Galstyan

Information Sciences Institute, University of Southern California

{ninarehm, umanggup, morstatt}@usc.edu, {gregv, galstyan}@isi.edu

Abstract

The widespread use of Artificial Intelligence (AI) in consequential domains, such as health-care and parole decision-making systems, has drawn intense scrutiny on the fairness of these methods. However, ensuring fairness is often insufficient as the rationale for a contentious decision needs to be audited, understood, and defended. We propose that the attention mechanism can be used to ensure fair outcomes while simultaneously providing feature attributions to account for how a decision was made. Toward this goal, we design an attention-based model that can be leveraged as an attribution framework. It can identify features responsible for both performance and fairness of the model through attention interventions and attention weight manipulation. Using this attribution framework, we then design a post-processing bias mitigation strategy and compare it with a suite of baselines. We demonstrate the versatility of our approach by conducting experiments on two distinct data types, tabular and textual.

1 Introduction

Machine learning algorithms that optimize for performance (e.g., accuracy) often result in unfair outcomes (Mehrabi et al., 2021). These algorithms capture biases present in the training datasets causing discrimination toward different groups. As machine learning continues to be adopted into fields where discriminatory treatments can lead to legal penalties, fairness and interpretability have become a necessity and a legal incentive in addition to an ethical responsibility (Barocas and Selbst, 2016; Hacker et al., 2020). Existing methods for fair machine learning include applying complex transformations to the data so that resulting representations are fair (Gupta et al., 2021; Moyer et al., 2018; Roy and Boddeti, 2019; Jaiswal et al., 2020; Song et al., 2019), adding regularizers to incorporate fairness (Zafar et al., 2017;

Kamishima et al., 2012; Mehrabi et al., 2020), or modifying the outcomes of unfair machine learning algorithms to ensure fairness (Hardt et al., 2016), among others. Here we present an alternative approach, which works by identifying the significance of different features in causing unfairness and reducing their effect on the outcomes using an attention-based mechanism.

With the advancement of transformer models and the attention mechanism (Vaswani et al., 2017), recent research in Natural Language Processing (NLP) has tried to analyze the effects and the interpretability of the attention weights on the decision making process (Wiegrefe and Pinter, 2019; Jain and Wallace, 2019; Serrano and Smith, 2019; Hao et al., 2021). Taking inspiration from these works, we propose to use an attention-based mechanism to study the fairness of a model. The attention mechanism provides an intuitive way to capture the effect of each attribute on the outcomes. Thus, by introducing the attention mechanism, we can analyze the effect of specific input features on the model’s fairness. We form visualizations that explain model outcomes and help us decide which attributes contribute to accuracy vs. fairness. We also show and confirm the observed effect of indirect discrimination in previous work (Zliobaite, 2015; Hajian and Domingo-Ferrer, 2013; Zhang et al., 2017) in which even with the absence of the sensitive attribute, we can still have an unfair model due to the existence of proxy attributes. Furthermore, we show that in certain scenarios those proxy attributes contribute more to the model unfairness than the sensitive attribute itself.

Based on the above observations, we propose a post-processing bias mitigation technique by diminishing the weights of features most responsible for causing unfairness. We perform studies on datasets with different modalities and show the flexibility of our framework on both tabu-

lar and large-scale text data, which is an advantage over existing interpretable non-neural and non-attention-based models. Furthermore, our approach provides a competitive and interpretable baseline compared to several recent fair learning techniques.

2 Approach

In this section, we describe our classification model that incorporates the attention mechanism. It can be applied to both text and tabular data and is inspired by works in attention-based models in text-classification (Zhou et al., 2016). We incorporate attention over the input features. Next, we describe how this attention over features can attribute the model’s unfairness to certain features. Finally, using this attribution framework, we propose a post-processing approach for mitigating unfairness.

In this work, we focus on binary classification tasks. We assume access to a dataset of triplets $\mathcal{D} = \{x_i, y_i, a_i\}_{i=1}^N$, where x_i, y_i, a_i are i.i.d. samples from data distribution $p(\mathbf{x}, \mathbf{y}, \mathbf{a})$. $\mathbf{a} \in \{a_1, \dots, a_l\}$ is a discrete variable with l possible values and denotes the sensitive or protected attributes with respect to which we want to be fair, $\mathbf{y} \in \{0, 1\}$ is the true label, $\mathbf{x} \in \mathbb{R}^m$ are features of the sample which may include sensitive attributes. We use \hat{y}_o to denote the binary outcome of the original model, and \hat{y}_z^k will represent the binary outcome of a model in which the attention weights corresponding to k^{th} feature are zeroed out. Our framework is flexible and general that it can be used to find attribution for any fairness notion. More particularly, we work with the group fairness measures like *Statistical Parity* (Dwork et al., 2012), *Equalized Odds* (Hardt et al., 2016), and *Equality of Opportunity* (Hardt et al., 2016), which are defined as:¹


Statistical Parity Difference (SPD):

$$\text{SPD}(\hat{\mathbf{y}}, \mathbf{a}) = \max_{a_i, a_j} |P(\hat{\mathbf{y}} = 1 \mid \mathbf{a} = a_i) - P(\hat{\mathbf{y}} = 1 \mid \mathbf{a} = a_j)|$$

Equality of Opportunity Difference (EqOpp):

$$\text{EqOpp}(\hat{\mathbf{y}}, \mathbf{a}, \mathbf{y}) = \max_{a_i, a_j} |P(\hat{\mathbf{y}} = 1 \mid \mathbf{a} = a_i, \mathbf{y} = 1) - P(\hat{\mathbf{y}} = 1 \mid \mathbf{a} = a_j, \mathbf{y} = 1)|$$

¹We describe and use the definition of these fairness measures as implemented in *Fairlearn* package (Bird et al., 2020).



(a) Classification model. (b) Attribution framework.

Figure 1: (a) In general classification model, for each feature f_k a vector representation e_k of length d^e is learned. This is passed to the attention layer which produces a d^e -dimensional vector representation for the sample instance i which is passed to two dense layers to get the final classification output. (b) The Attribution framework has the same architecture as the general model. One outcome is obtained through the original model and another through the model that has some attention weights zeroed. The observed difference in accuracy and fairness measures will indicate the effect of the zeroed out features on accuracy and fairness.

Equalized Odds Difference (EqOdd):

$$\text{EqOdd}(\hat{\mathbf{y}}, \mathbf{a}, \mathbf{y}) = \max_{a_i, a_j} \max_{y \in \{0, 1\}} |P(\hat{\mathbf{y}} = 1 \mid \mathbf{a} = a_i, \mathbf{y} = y) - P(\hat{\mathbf{y}} = 1 \mid \mathbf{a} = a_j, \mathbf{y} = y)|$$

2.1 General Model: Incorporating Attention over Inputs in Classifiers

We consider each feature value as an individual entity (like the words are considered in text-classification) and learn a fixed-size embedding $\{e_k\}_{k=1}^m$, $e_k \in \mathbb{R}^{d^e}$ for each feature, $\{f_k\}_{k=1}^m$. These vectors are passed to the attention layer. The Computation of attention weights and the final representation for a sample is described in Eq. 1. $E = [e_1 \dots e_m]$, $E \in \mathbb{R}^{d^e \times m}$ is the concatenation of all the embeddings, $w \in \mathbb{R}^{d^e}$ is a learnable parameter, $r \in \mathbb{R}^{d^e}$ denotes the overall sample representation, and $\alpha \in \mathbb{R}^m$ denotes the attention weights.

$$H = \tanh(E); \alpha = \text{softmax}(w^T H); r = \tanh(E\alpha^T) \quad (1)$$

The resulting representation, r , is passed to the feed-forward layers for classification. In this work, we have used two feed-forward layers (See Fig. 1 for overall architecture).

2.2 Fairness Attribution with Attention

The aforementioned classification model with the attention mechanism combines input feature em-

Algorithm 1: Bias Mitigation by Attention

```
1 Input: decay rate  $d_r$  ( $0 \leq d_r < 1$ ),  $n$  test
  samples indexed by variable  $i$ .
2 Output: final predictions, unfair features.
3 Calculate the attention weights  $\alpha_{ki}$  for  $k^{\text{th}}$ 
  feature in sample  $i$  using the attention
  layer as in Eq. 1.
4 unfair_feature_set = { }
5 for each feature (index)  $k$  do
6   if  $SPD(\hat{y}_o, \mathbf{a}) - SPD(\hat{y}_z^k, \mathbf{a}) \geq 0$  then
7     unfair_feature_set =
8     | unfair_feature_set  $\cup \{k\}$ 
9   end
10 end
11 for each feature (index)  $k$  do
12   if  $k$  in unfair_feature_set then
13     Set  $\alpha_{ki} \leftarrow (d_r \times \alpha_{ki})$  for all  $n$ 
14     samples
15   end
16 end
17 Use new attention weights to obtain the
  final predictions  $\hat{Y}$ .
18 return  $\hat{Y}$ , unfair_feature_set
```

beddings by taking a weighted combination. By manipulating the weights, we can intuitively capture the effects of specific features on the output. To this end, we observe the effect of each attribute on the fairness of outcomes by zeroing out or reducing its attention weights and recording the change. Other works have used similar ideas to understand the effect of attention weights on accuracy and evaluate interpretability of the attention weights by comparing the difference in outcomes in terms of measures such as Jensen-Shannon Divergence (Serrano and Smith, 2019) but not for fairness. We are interested in the effect of features on fairness measures. Thus, we measure the difference in fairness of the outcomes based on the desired fairness measure. A large change in fairness measure and a small change in performance of the model would indicate that this feature is mostly responsible for unfairness, and it can be dropped without causing large impacts on performance. The overall framework is shown in Fig. 1. First, the outcomes are recorded with the original attention weights intact (Fig. 1a). Next, attention weights corresponding to a particular feature are zeroed out, and the difference in performance and fairness measures is recorded (Fig. 1b).

Based on the observed differences, one may conclude how incorporating this feature contributes to fairness/unfairness.

To measure the effect of the k^{th} feature on different fairness measures, we consider the difference in the fairness of outcomes of the original model and model with k^{th} feature’s effect removed. For example, for statistical parity difference, we will consider $SPD(\hat{y}_o, \mathbf{a}) - SPD(\hat{y}_z^k, \mathbf{a})$. A negative value will indicate that the k^{th} feature helps mitigate unfairness, and a positive value will indicate that the k^{th} feature contributes to unfairness. This is because \hat{y}_z^k captures the exclusion of the k^{th} feature (zeroed out attention weight for that feature) from the decision-making process. If the value is positive, it indicates that not having this feature makes the bias lower than when we include it. Notice here, we focus on global attribution, so we measure this over all the samples; however, this can also be turned into local attribution by focusing on individual sample i only.

2.3 Bias Mitigation by Removing Unfair Features

As discussed in the previous section, we can identify features that contribute to unfair outcomes according to different fairness measures. A simple technique to mitigate or reduce bias is to reduce the attention weights of these features. This mitigation technique is outlined in Algorithm 1. In this algorithm, we first individually set attention weights for each of the features in all the samples to zero and monitor the effect on the desired fairness measure. We have demonstrated the algorithm for SPD, but other measures, such as EqOdd, EqOpp, and even accuracy can be used (in which case the “unfair_feature_set” can be re-named to feature set which harms accuracy instead of fairness). If the k^{th} feature contributes to unfairness, we reduce its attention weight using decay rate value. This is because \hat{y}_z^k captures the exclusion of the k^{th} feature (zeroed attention weight for that feature) compared to the original outcome \hat{y}_o for when all the feature weights are intact; otherwise, we use the original attention weight. We can also control the fairness-accuracy trade-off by putting more attention weight on features that boost accuracy while keeping the fairness of the model the same and down-weighting features that hurt accuracy, fairness, or both.

This post-processing technique has a couple of

advantages over previous works in bias mitigation or fair classification approaches. First, the post-processing approach is computationally efficient as it does not require model retraining to ensure fairness for each sensitive attribute separately. Instead, the model is trained once by incorporating all the attributes, and then one manipulates attention weights during test time according to particular needs and use-cases. Second, the proposed mitigation method provides an explanation and can control the fairness-accuracy trade-off. This is because manipulating the attention weights reveals which features are important for getting the desired outcome, and by how much. This provides an explanation for the outcome and also a mechanism to control the fairness-accuracy trade-off by the amount of the manipulation.

3 Experimental Setup

We perform a suite of experiments on synthetic and real-world datasets to evaluate our attention based interpretable fairness framework. The experiments on synthetic data are intended to elucidate interpretability in controlled settings, where we can manipulate the relations between input and output feature. The experiments on real-world data aim to validate the effectiveness of the proposed approach on both tabular and non-tabular (textual) data.

3.1 Types of Experiments

We enumerate the experiments and their goals as follows:

Experiment 1: Attributing Fairness with Attention The purpose of this experiment is to demonstrate that our attribution framework can capture correct attributions of features to fairness outcomes. We present our results for tabular data in Sec. 4.1.

Experiment 2: Bias Mitigation via Attention Weight Manipulation In this experiment, we seek to validate the proposed post-processing bias mitigation framework and compare it with various recent mitigation approaches. The results for real-world tabular data are presented in Sec. 4.2.

Experiment 3: Validation on Textual Data The goal of this experiment is to demonstrate the flexibility of the proposed attention-based method by conducting experiments on non-tabular, textual data. The results are presented in Sec. 4.3.

3.2 Datasets

3.2.1 Synthetic Data

To validate the attribution framework, we created two synthetic datasets in which we control how features interact with each other and contribute to the accuracy and fairness of the outcome variable. These datasets capture some of the common scenarios, namely the data imbalance (skewness) and indirect discrimination issues, arising in fair decision or classification problems.

Scenario 1: First, we create a simple scenario to demonstrate that our framework identifies correct feature attributions for fairness and accuracy. We create a feature that is correlated with the outcome (responsible for accuracy), a discrete feature that causes the prediction outcomes to be biased (responsible for fairness), and a continuous feature that is independent of the label or the task (irrelevant for the task). For intuition, suppose the attention-based attribution framework works correctly. In this case, we expect to see a reduction in accuracy upon removing (i.e., making the attention weight zero) the feature responsible for the accuracy, reduction in bias upon removing the feature responsible for bias, and very little or no change upon removing the irrelevant feature. With this objective, we generated a synthetic dataset with three features, i.e., $x = [f_1, f_2, f_3]$ as follows²:

$$f_1 \sim \text{Ber}(0.9) \quad f_2 \sim \text{Ber}(0.5) \quad f_3 \sim \mathcal{N}(0, 1)$$

$$y \sim \begin{cases} \text{Ber}(0.9) & \text{if } f_2 = 1 \\ \text{Ber}(0.1) & \text{if } f_2 = 0 \end{cases}$$

Clearly, f_2 has the most predictive information for the task and is responsible for accuracy. Here, we consider f_1 as the sensitive attribute. f_1 is an imbalanced feature that can bias the outcome and is generated such that there is no intentional correlation between f_1 and the outcome, y or f_2 . f_3 is sampled from a normal distribution independent of the outcome y , or the other features, making it irrelevant for the task. Thus, an ideal classifier would be fair if it captures the correct outcome without being affected by the imbalance in f_1 . However, due to limited data and skew in f_1 , there will be some undesired bias — few errors when $f_1 = 0$ can lead to large statistical parity.

²We use $x \sim \text{Ber}(p)$ to denote that x is a Bernoulli random variable with $P(x = 1) = p$.

Scenario 2: Using features that are not identified as sensitive attributes can result in unfair decisions due to their implicit relations or correlations with the sensitive attributes. This phenomenon is called indirect discrimination (Zliobaite, 2015; Hajian and Domingo-Ferrer, 2013; Zhang et al., 2017). We designed this synthetic dataset to demonstrate and characterize the behavior of our framework under indirect discrimination. Similar to the previous scenario, we consider three features. Here, f_1 is considered as the sensitive attribute, and f_2 is correlated with f_1 and the outcome, y . The generative process is as follows:

$$f_1 \sim \begin{cases} \text{Ber}(0.9) & \text{if } f_2 = 1 \\ \text{Ber}(0.1) & \text{if } f_2 = 0 \end{cases} \quad f_2 \sim \text{Ber}(0.5)$$

$$f_3 \sim \mathcal{N}(0, 1) \quad y \sim \begin{cases} \text{Ber}(0.7) & \text{if } f_2 = 1 \\ \text{Ber}(0.3) & \text{if } f_2 = 0 \end{cases}$$

In this case f_1 and y are correlated with f_2 . The model should mostly rely on f_2 for its decisions. However, due to the correlation between f_1 and f_2 , we expect f_2 to affect both the accuracy and fairness of the model. Thus, in this case, indirect discrimination is possible. Using such a synthetic dataset, we demonstrate a) indirect discrimination and b) the need to have an attribution framework to reason about unfairness and not blindly focus on the sensitive attributes for bias mitigation.

3.2.2 Real-world Datasets

We demonstrate our approach on the following real-world datasets:

Tabular Datasets: We conduct our experiments on two real-world tabular datasets — *UCI Adult* (Dua and Graff, 2017) and *Heritage Health*³ datasets. The *UCI Adult* dataset contains census information about individuals, with the prediction task being whether the income of the individual is higher than \$50k or not. The sensitive attribute, in this case, is gender (male/female). The *Heritage Health* dataset contains patient information, and the task is to predict the Charlson Index (comorbidity index, which is a patient survival indicator). Each patient is grouped into one of the 9 possible age groups, and we consider this as the sensitive attribute. We used the same pre-processing and train-test splits as in Gupta et al. (2021).

Non-Tabular or Text Dataset: We also experiment with a non-tabular, text dataset. We used

the *biosbias* dataset (De-Arteaga et al., 2019). The dataset contains short bios of individuals. The task is to predict the occupation of the individual from their bio. We utilized the bios from the year 2018 from the 2018_34 archive and considered two occupations for our experiments, namely, nurse and dentist. The dataset was split into 70-15-15 train, validation, and test splits. De-Arteaga et al. (2019) has demonstrated the existence of gender bias in this prediction task and showed that certain gender words are associated with certain job types (e.g., *she* to nurse and *he* to dentist).

3.3 Bias Mitigation Baselines

For our baselines, we consider methods that learn representations of data so that information about sensitive attributes is eliminated. *CVIB* (Moyer et al., 2018) realizes this objective through a conditional variational autoencoder, whereas *MIFR* (Song et al., 2019) uses a combination of information bottleneck term and adversarial learning to optimize the fairness objective. *FCRL* (Gupta et al., 2021) optimizes information theoretic objectives that can be used to achieve good trade-offs between fairness and accuracy by using specialized contrastive information estimators. In addition to information-theoretic approaches, we also considered baselines that use adversarial learning such as *MaxEnt-ARL* (Roy and Boddeti, 2019), *LAFTR* (Madras et al., 2018), and *Adversarial Forgetting* (Jaiswal et al., 2020). Note that in contrast to our approach, the baselines described above are not interpretable as they are incapable of directly attributing features to fairness outcomes. For the textual data, we compare our approach with the debiasing technique proposed in De-Arteaga et al. (2019), which works by masking the gender-related words and then training the model on this masked data.

4 Results

4.1 Attributing Fairness with Attention

First, we test our method’s ability to capture correct attributions in controlled experiments with synthetic data (described in Sec. 3.2.1). We also conduct a similar experiment with *UCI Adult* and *Heritage Health* datasets which can be found in the appendix. Fig. 2 summarizes our results by visualizing the attributions, which we now discuss.

In *Scenario 1*, as expected, f_2 is correctly attributed to being responsible for the accuracy and

³<https://www.kaggle.com/c/hhp>

removing it hurts the accuracy drastically. Similarly, f_1 is correctly shown to be responsible for unfairness and removing it creates a fairer outcome. Ideally, the model should not be using any information about f_1 as it is independent of the task, but it does. Therefore, by removing f_1 , we can ensure that information is not used and hence outcomes are fair. Lastly, as expected, f_3 was the irrelevant feature, and its effects on accuracy and fairness are negligible.

In *Scenario 2*, our framework captures the effect of indirect discrimination. We can see that removing f_2 reduces bias as well as accuracy drastically. This is because f_2 is the predictive feature, but due to its correlation with f_1 , it can also indirectly affect the model’s fairness. More interestingly, although f_1 is the sensitive feature, removing it does not play a drastic role in fairness or the accuracy. This is an important finding as it shows why removing f_1 on its own can not give us a fairer model due to the existence of correlations to other features and indirect discrimination. Overall, our results are intuitive and thus validate our assumption that attention-based framework can provide reliable feature attributions for the fairness and accuracy of the model.

4.2 Attention as a Mitigation Technique

As we have highlighted earlier, understanding how the information within features interact and contribute to the decision making can be used to design effective bias mitigation strategies. One such example was shown in Sec. 4.1. Often real-world datasets have features which cause indirect discrimination, due to which fairness can not be achieved by simply eliminating the sensitive feature from the decision process. Using the attributions derived from our attention-based attribution framework, we propose a post-processing mitigation strategy. Our strategy is to intervene on attention weights as discussed in Sec. 2.3. We first attribute and identify the features responsible for the unfairness of the outcomes, i.e., all the features whose exclusion will decrease the bias compared to the original model’s outcomes and gradually decrease their attention weights to zero as also outlined in Algorithm 1. We do this by first using the whole fraction of the attention weights learned and gradually use less fraction of the weights until the weights are completely zeroed out.

For all the baselines described in Sec. 3.3, we

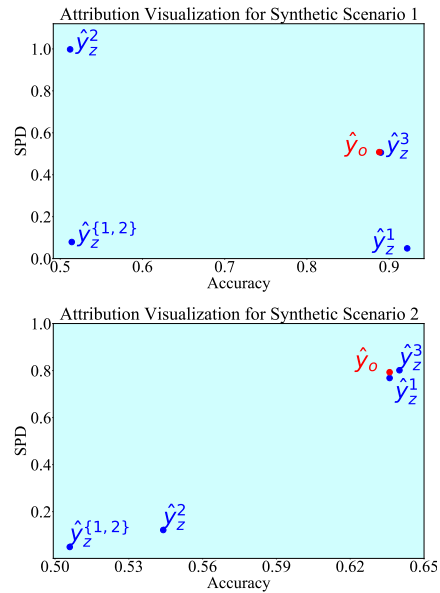


Figure 2: Results from the synthetic datasets. Following the \hat{y}_o and \hat{y}_z notations, \hat{y}_o represents the original model outcome with all the attention weights intact, while \hat{y}_z^k represents the outcome of the model in which the attention weights corresponding to k^{th} feature are zeroed out (e.g. \hat{y}_z^1 represents when attention weights of feature f_1 are zeroed out). The results show how the accuracy and fairness (SPD) of the model change by exclusion of each feature.

used the approach outlined in Gupta et al. (2021) for training a downstream classifier and evaluating the accuracy/fairness trade-offs. The downstream classifier was a 1-hidden-layer MLP with 50 neurons along with ReLU activation function. Each method was trained with five different seeds, and we report the average accuracy and fairness measure as statistical parity difference (SPD). Results for other fairness notions can be found in the appendix. *CVIB*, *MaxEnt-ARL*, *Adversarial Forgetting* and *FCRL* are designed for statistical parity notion of fairness and are not applicable for other measures like Equalized Odds and Equality of Opportunity. *LAFTR* can only deal with binary sensitive attributes and thus not applicable for Heritage Health dataset. Notice that our approach does not have these limitations. For our approach, we vary the attention weights and report the resulting fairness-accuracy trade offs.

Fig. 3 compares fairness-accuracy trade-offs of different bias mitigation approaches. We desire outcomes to be fairer, i.e., lower values of SPD and to be more accurate, i.e., towards the right. The results show that using attention attributions can indeed be beneficial for reducing bias. Moreover, our mitigation framework based on the ma-




Figure 3: Accuracy vs parity curves for UCI Adult and Heritage Health datasets.

nipulation of the attention weights is competitive with state-of-the-art mitigation strategies. However, most of these approaches are specifically designed and optimized to achieve parity and do not provide any interpretability. Our model can not only achieve comparable and competitive results, but it is also able to provide explanation such that the users exactly know what feature and by how much it was manipulated to get the corresponding outcome. Another advantage of our model is that it needs only one round of training. The adjustments to attention weights are made post-training; thus, it is possible to achieve different trade-offs. Moreover, our approach does not need to know sensitive attributes while training; thus, it could work with other sensitive attributes not known beforehand or during training. Lastly, here we merely focused on mitigating bias as our goal was to show that the attribution framework can identify problematic features and their removal would result in bias mitigation. We manipulated attention weights of all the features that contributed to unfairness irrespective of if they helped maintaining high accuracy or not. However, the trade-off results can be improved by carefully considering the trade-off each feature contributes to with regards to both accuracy and fairness to achieve better trade-off results which can be investigated as a future direction. The advantage of our work is that this trade-off curve can be controlled by controlling how many features and by how much to be manipulated.

4.3 Experiments with Non-Tabular Data

In addition to providing interpretability, our approach is flexible and useful for controlling fairness in modalities other than tabular datasets. To

put this to the test, we applied our model to mitigate bias in text-based data. We consider the *bios-bias* dataset (De-Arteaga et al., 2019), and use our mitigation technique to reduce observed biases in the classification task performed on this dataset. We compare our approach with the debiasing technique proposed in the original paper (De-Arteaga et al., 2019), which works by masking the gender-related words and then training the model on this masked data. As discussed earlier, such a method is computationally inefficient. It requires re-training the model or creating a new masked dataset, each time it is required to debias the model against different attributes, such as gender vs. race. For the baseline pre-processing method, we masked the gender-related words, such as names and gender words, as provided in the *biosbias* dataset and trained the model on the filtered dataset. On the other hand, we trained the model on the raw bios for our post-processing method and only manipulated attention weights of the gender words during the testing process as also provided in the *biosbias* dataset. In order to measure the bias, we used the same measure as in (De-Arteaga et al., 2019) which is based on the equality of opportunity notion of fairness (Hardt et al., 2016) and reported the True Positive Rate Difference (TPRD) for each occupation amongst different genders. As shown in Table 1, our post-processing mitigation technique provides lower TRPD while being more accurate, followed by the technique that masks the gendered words before training. Although both methods reduce the bias compared to a model trained on raw bios without applying any mask or invariance to gendered words, our post-processing method

Method	Dentist TPRD (stdev)	Nurse TPRD (stdev)	Accuracy (stdev)
Post-Processing (Ours)	0.0202 (0.010)	0.0251 (0.020)	0.951 (0.013)
Pre-Processing	0.0380 (0.016)	0.0616 (0.025)	0.946 (0.011)
Not Debaised Model	0.0474 (0.025)	0.1905 (0.059)	0.958 (0.011)

Table 1: Difference of the True Positive Rates (TPRD) amongst different genders for the dentist and nurse occupations on the biosbias dataset. Our introduced post-processing method is the most effective in reducing the disparity for both occupations compared to the pre-processing technique.

Post-Processing (Ours)	Pre-Processing	Not Debaised Model
She practices in Apo, Armed Forces Europe and has the professional credentials of R.N. The NPI Number for Rebekah Bushey is 1073935136 and she holds a License No. RN.0172354 (Colorado).	She practices in Apo, Armed Forces Europe and has the professional credentials of R.N.. The NPI Number for Rebekah Bushey is 1073935136 and she holds a License No. RN.0172354 (Colorado).	She practices in Apo, Armed Forces Europe and has the professional credentials of R.N.. The NPI Number for Rebekah Bushey is 1073935136 and she holds a License No. RN.0172354 (Colorado).
Post-Processing (Ours)	Pre-Processing	Not Debaised Model
She has worked inpatient and outpatient from pediatrics to adults. She is currently working on obtaining her Doctorate of Nursing specializing in psychiatry and mental health .	She has worked inpatient and outpatient from pediatrics to adults. She is currently working on obtaining her Doctorate of Nursing specializing in psychiatry and mental health.	She has worked inpatient and outpatient from pediatrics to adults. She is currently working on obtaining her Doctorate of Nursing specializing in psychiatry and mental health .

Figure 4: Qualitative results from the non-tabular data experiment on the job classification task based on bio texts. Green regions are the top three words used by the model for its prediction based on the attention weights. While the Not Debaised Model mostly focuses on gendered words, our method focused on profession-based words, such as R.N. (Registered Nurse), to correctly predict “nurse.”

is more effective. Fig. 4 also highlights qualitative differences between models in terms of their most attentive features for the prediction task. As shown in the results, our post-processing technique is able to use more meaningful words, such as R.N. (registered nurse) to predict the outcome label nurse compared to both baselines, while the non-debaised model focuses on gendered words.

5 Related Work

Fairness. The research in fairness concerns itself with various topics (Mehrabi et al., 2021). In this work, we utilized different metrics that were introduced previously (Dwork et al., 2012; Hardt et al., 2016), to measure the amount of bias. We also used different bias mitigation strategies to compare against our mitigation strategy, such as FCRL (Gupta et al., 2021), CVIB (Moyer et al., 2018), MIFR (Song et al., 2019), adversarial forgetting (Jaiswal et al., 2020), MaxEnt-ARL (Roy and Boddeti, 2019), and LAFTR (Madras et al., 2018). We also utilized concepts and datasets that were analyzing existing biases in NLP systems, such as (De-Arteaga et al., 2019) which studied the existing biases in NLP systems on the occupation classification task on the bios dataset.

Interpretability. There is a body of work in NLP literature that tried to analyze the effect of the attention weights on interpretability of the model (Wiegrefe and Pinter, 2019; Jain and Wallace, 2019; Serrano and Smith, 2019). Other work

also utilized attention weights to define an attribution score to be able to reason about how transformer models such as BERT work (Hao et al., 2021). Notice that although Jain and Wallace (2019) claim that attention might not be explanation, a body of work has proved otherwise including (Wiegrefe and Pinter, 2019) in which authors directly target the work in Jain and Wallace (2019) and analyze in detail the problems associated with this study. In addition, Vig et al. (2020) analyze the effect of the attention weights in transformer models for bias analysis in language models.

6 Discussion

In this work, we analyzed how attention weights contribute to fairness and accuracy of a predictive model. We proposed an attribution method that leverages the attention mechanism and showed the effectiveness of this approach on both tabular and text data. Using this interpretable attribution framework we then introduced a post-processing bias mitigation strategy based on attention weight manipulation. We validated the proposed framework by conducting experiments with different baselines, fairness metrics, and data modalities.

Acknowledgments

We thank anonymous reviewers for providing insightful feedback. Ninareh Mehrabi’s research was funded by USC + Amazon Center on Secure and Trusted Machine Learning fellowship.

Broader Impact

Although our work can have a positive impact in allowing to reason about fairness and accuracy of models and reduce their bias, it can also have negative societal consequences if used unethically. For instance, it has been previously shown that interpretability frameworks can be used as a means for fairwashing which is when malicious users generate fake explanations for their unfair decisions to justify them (Anders et al., 2020). In addition, previously it has been shown that interpretability frameworks are vulnerable against adversarial attacks (Slack et al., 2020). We acknowledge that our framework may also be targeted by malicious users for malicious intent that can manipulate attention weights to either generate fake explanations or unfair outcomes. We also acknowledge that our method is not achieving the best accuracy-fairness trade-off on the UCI Adult dataset for the statistical parity notion of fairness and has room for improvement.

References

- Christopher Anders, Plamen Pasliev, Ann-Kathrin Dombrowski, Klaus-Robert Müller, and Pan Kessel. 2020. [Fairwashing explanations with off-manifold detergent](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 314–323. PMLR.
- Solon Barocas and Andrew D Selbst. 2016. Big data’s disparate impact. *Calif. L. Rev.*, 104:671.
- Sarah Bird, Miro Dudík, Richard Edgar, Brandon Horn, Roman Lutz, Vanessa Milan, Mehrnoosh Sameki, Hanna Wallach, and Kathleen Walker. 2020. [Fairlearn: A toolkit for assessing and improving fairness in ai](#). Technical Report MSR-TR-2020-32, Microsoft.
- Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. 2019. Bias in bios: A case study of semantic representation bias in a high-stakes setting. In *proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 120–128.
- Dheeru Dua and Casey Graff. 2017. [UCI machine learning repository](#).
- Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. [Fairness through awareness](#). In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS ’12, pages 214–226, New York, NY, USA. ACM.
- Umang Gupta, Aaron M Ferber, Bistra Dilkina, and Greg Ver Steeg. 2021. Controllable guarantees for fair outcomes via contrastive information estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7610–7619.
- Philipp Hacker, Ralf Krestel, Stefan Grundmann, and Felix Naumann. 2020. Explainable ai under contract and tort law: legal incentives and technical challenges. *Artificial Intelligence and Law*, pages 1–25.
- S. Hajian and J. Domingo-Ferrer. 2013. [A methodology for direct and indirect discrimination prevention in data mining](#). *IEEE Transactions on Knowledge and Data Engineering*, 25(7):1445–1459.
- Yaru Hao, Li Dong, Furu Wei, and Ke Xu. 2021. Self-attention attribution: Interpreting information interactions inside transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12963–12971.
- Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. 2016. [Equality of opportunity in supervised learning](#). In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Sarthak Jain and Byron C. Wallace. 2019. [Attention is not Explanation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3543–3556, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ayush Jaiswal, Daniel Moyer, Greg Ver Steeg, Wael AbdAlmageed, and Premkumar Natarajan. 2020. [Invariant representations through adversarial forgetting](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):4272–4279.
- Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. 2012. Fairness-aware classifier with prejudice remover regularizer. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 35–50. Springer.
- David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. 2018. [Learning adversarially fair and transferable representations](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3384–3393. PMLR.
- Ninareh Mehrabi, Yuzhong Huang, and Fred Morstatter. 2020. Statistical equity: A fairness classification objective. *arXiv preprint arXiv:2005.07293*.

- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. [A survey on bias and fairness in machine learning](#). *ACM Comput. Surv.*, 54(6).
- Daniel Moyer, Shuyang Gao, Rob Brekelmans, Aram Galstyan, and Greg Ver Steeg. 2018. [Invariant representations without adversarial training](#). In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Proteek Chandan Roy and Vishnu Naresh Boddeti. 2019. Mitigating information leakage in image representations: A maximum entropy approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2586–2594.
- Sofia Serrano and Noah A. Smith. 2019. [Is attention interpretable?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2931–2951, Florence, Italy. Association for Computational Linguistics.
- Dylan Slack, Sophie Hilgard, Emily Jia, Sameer Singh, and Himabindu Lakkaraju. 2020. Fooling lime and shap: Adversarial attacks on post hoc explanation methods. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 180–186.
- Jiaming Song, Pratyusha Kalluri, Aditya Grover, Shengjia Zhao, and Stefano Ermon. 2019. [Learning controllable fair representations](#). In *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 2164–2173. PMLR.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. 2020. [Investigating gender bias in language models using causal mediation analysis](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 12388–12401. Curran Associates, Inc.
- Sarah Wiegrefe and Yuval Pinter. 2019. [Attention is not not explanation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 11–20, Hong Kong, China. Association for Computational Linguistics.
- Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rogriguez, and Krishna P. Gummadi. 2017. [Fairness Constraints: Mechanisms for Fair Classification](#). In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 962–970, Fort Lauderdale, FL, USA. PMLR.
- Lu Zhang, Yongkai Wu, and Xintao Wu. 2017. [A causal framework for discovering and removing direct and indirect discrimination](#). In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 3929–3935.
- Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. 2016. [Attention-based bidirectional long short-term memory networks for relation classification](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 207–212, Berlin, Germany. Association for Computational Linguistics.
- Indre Zliobaite. 2015. A survey on measuring indirect discrimination in machine learning. *arXiv preprint arXiv:1511.00148*.

A Appendix

We included additional bias mitigation results using other fairness metrics, such as equality of opportunity and equalized odds on both of the Adult and Heritage Health datasets in this supplementary material. We also included additional post-processing results along with additional qualitative results both for the tabular and non-tabular dataset experiments. More details can be found under each sub-section.

A.1 Results on Tabular Data

Here, we show the results of our mitigation framework considering equality of opportunity and equalized odds notions of fairness. We included baselines that were applicable for these notions. Notice not all the baselines we used in our previous analysis for statistical parity were applicable for equality of opportunity and equalized odds notions of fairness; thus, we only included the applicable ones. In addition, LAFTR is only applicable when the sensitive attribute is a binary variable, so it was not applicable to be included in the analysis for the heritage health data where the sensitive attribute is non-binary. Results of these analysis is shown in Figures 8 and 9. We once again show competitive and comparable results to other baseline methods, while having the advantage of being interpretable and not requiring multiple trainings to satisfy different fairness notions or fairness on different sensitive attributes. Our framework is also flexible for different fairness measures and can be applied to binary or non-binary sensitive features.

In addition, we show how different features contribute differently under different fairness notions. Fig. 5 demonstrates the top three features that contribute to unfairness the most along with the percentages of the fairness improvement upon their removal for each of the fairness notions. As observed from the results, while equality of opportunity and equalized odds are similar in terms of their problematic features, statistical parity has different trends. This is also expected as equality of opportunity and equalized odds are similar fairness notions in nature compared to statistical parity.

We also compared our mitigation strategy with the Hardt et al. post-processing approach (Hardt et al., 2016). Using this post-processing imple-

Abbreviation	Meaning
PlaceSvcs	Place where the member was treated.
LOS	Length of stay.
dsfs	Days since first service that year.

Table 2: Some abbreviations used in Heritage Health dataset’s feature names. These abbreviations are listed for clarity of interpreting each feature’s meaning specifically in our qualitative analysis or attribution visualizations.

mentation⁴, we obtained the optimal solution that tries to satisfy different fairness notions subject to accuracy constraints. For our results, we put the results from zeroing out all the attention weights corresponding to the problematic features that were detected from our interpretability framework. However, notice that since our mitigation strategy can control different trade-offs we can have different results depending on the scenario. Here, we reported the results from zeroing out the problematic attention weights that is targeting fairness mostly. From the results demonstrated in Tables 3 and 4, we can see comparable numbers to those obtained from (Hardt et al., 2016). This again shows that our interpretability framework yet again captures the correct responsible features and that the mitigation strategy works as expected.

A.2 Results on non-tabular Data

We also included some additional qualitative results from the experiments on non-tabular data in Fig. 6.

A.3 Interpreting Fairness with Attention

Fig. 7 shows results on a subset of the features from the *UCI Adult* and *Heritage Health* datasets (to keep the plots uncluttered and readable, we incorporated the most interesting features in the plot), and provide some intuition about how different features in these datasets contribute to the model fairness and accuracy. While features such as *capital gain* and *capital loss* in the *UCI Adult* dataset are responsible for improving accuracy and reducing bias, we can observe that features such as *relationship* or *marital status*, which can be indirectly correlated with the feature *sex*, have a negative impact on fairness. For the *Heritage Health* dataset, including the features *drugCount ave* and *dsfs max* provide accuracy gains but at the expense of fairness, while including *no Claims* and *no Spe-*

⁴<https://fairlearn.org>




Figure 5: Top three features for each fairness definition removing which caused the most benefit in improving the corresponding fairness definition. The percentage of improvement upon removal is marked on the y -axis for adult and heritage health datasets.

Post-Processing (Ours)	Pre-Processing	Not Debiased Model
Megan was born and raised in Arlington, Texas. She graduated from Baylor University Louise Herrington School of Nursing in 2015 with her Bachelors of Science in Nursing . Following graduation, she completed a cardiac telemetry internship at Texas Health Resources Fort Worth. She continued to work there for two years taking care of postsurgical cardiac patients and taking an active role in patient education committees. Since that time, she has left the hospital setting and has joined the practice . Megan is recently married and has a rescue Chihuahua Terrier mix pup. In her free time, she loves spending time with family and friends, trying new foods, and catching up on a good book.	Megan was born and raised in Arlington, Texas. She graduated from Baylor University Louise Herrington School of Nursing in 2015 with her Bachelors of Science in Nursing. Following graduation, she completed a cardiac telemetry internship at Texas Health Resources Fort Worth. She continued to work there for two years taking care of postsurgical cardiac patients and taking an active role in patient education committees. Since that time, she has left the hospital setting and has joined the practice. Megan is recently married and has a rescue Chihuahua Terrier mix pup. In her free time, she loves spending time with family and friends, trying new foods, and catching up on a good book.	Megan was born and raised in Arlington, Texas. She graduated from Baylor University Louise Herrington School of Nursing in 2015 with her Bachelors of Science in Nursing . Following graduation, she completed a cardiac telemetry internship at Texas Health Resources Fort Worth. She continued to work there for two years taking care of postsurgical cardiac patients and taking an active role in patient education committees. Since that time, she has left the hospital setting and has joined the practice. Megan is recently married and has a rescue Chihuahua Terrier mix pup. In her free time, she loves spending time with family and friends, trying new foods, and catching up on a good book.

Figure 6: Additional qualitative results from the non-tabular data experiment on the job classification task based on the bio texts. Green regions represent top three words that the model used for its prediction based on the attention weights.

cialities negatively impact both accuracy and fairness.

A.4 Information on Datasets and Features

More details about each of the datasets along with the descriptions of each feature for the Adult dataset can be found at⁵ and for the Heritage Health dataset can be found at⁶. In our qualitative results, we used the feature names as marked in these datasets. If the names or acronyms are unclear kindly reference to the references mentioned for more detailed description for each of the features. Although most of the features in the Adult datasets are self-descriptive, Heritage Health dataset includes some abbreviations that we list in Table 2 for the ease of interpreting each feature’s meaning.

⁵<https://archive.ics.uci.edu/ml/datasets/adult>

⁶<https://www.kaggle.com/c/hhp>




Figure 7: Results from the real-world datasets. Note that in our \hat{y}_z notation we replaced indexes with actual feature names for clarity in these results on real-world datasets as there is not one universal indexing schema, but the feature names are more universal and descriptive for this case. Labels on the points represent the feature name that was removed (zeroed out) according to our \hat{y}_z notation. The results show how the accuracy and fairness of the model (in terms of statistical parity difference) change by exclusion of each feature.

	Accuracy	SPD	Accuracy	EQOP	Accuracy	EQOD
Attention (Ours)	0.77 (0.006)	0.012 (0.003)	0.81 (0.013)	0.020 (0.019)	0.81 (0.021)	0.027 (0.023)
Hardt et al.	0.77 (0.012)	0.013 (0.005)	0.83 (0.005)	0.064 (0.016)	0.81 (0.007)	0.047 (0.014)

Table 3: Adult results on post-processing approach from Hardt et al. vs our attention method when all problematic features are zeroed out.

	Accuracy	SPD	Accuracy	EQOP	Accuracy	EQOD
Attention (Ours)	0.68 (0.004)	0.04 (0.015)	0.68 (0.015)	0.15 (0.085)	0.68 (0.015)	0.10 (0.085)
Hardt et al.	0.68 (0.005)	0.05 (0.018)	0.75 (0.001)	0.20 (0.033)	0.69 (0.012)	0.19 (0.031)

Table 4: Heritage Health results on post-processing approach from Hardt et al. vs our attention method when all problematic features are zeroed out.




Figure 8: Accuracy vs equality of opportunity curves for UCI Adult and Heritage Health datasets.

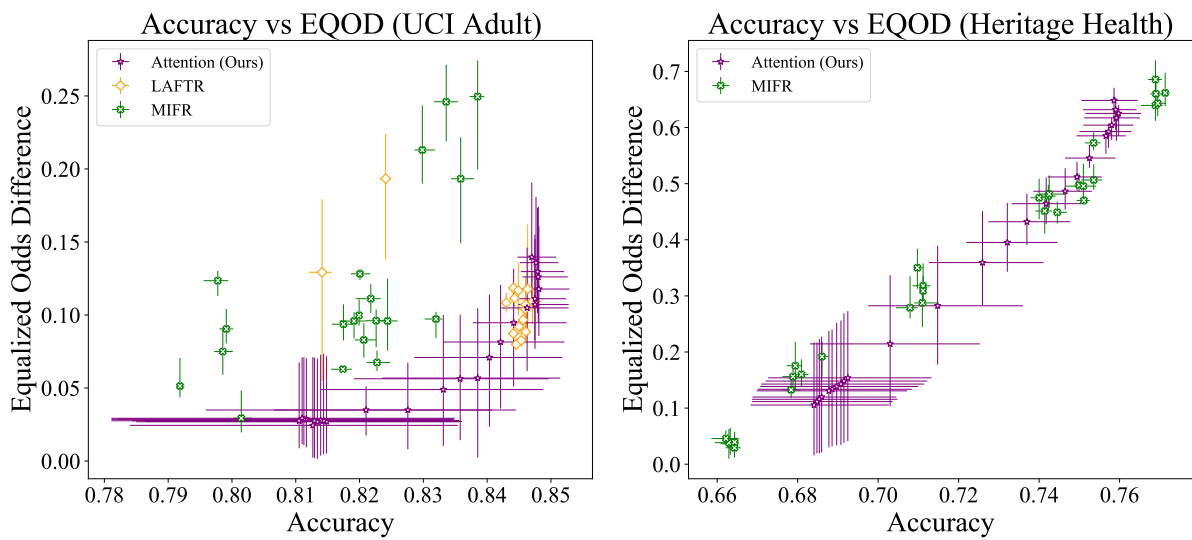


Figure 9: Accuracy vs equalized odds curves for UCI Adult and Heritage Health datasets.